Building Code with Dynamic Staging

Piotr Danilewski[1,2,3] Philipp Slusallek[1,2,4]

[1]Saarland University, Germany [2]Intel Visual Computing Institute, Germany
[3]Theoretical Computer Science, Jagiellonian University, Poland
[4]Deutsches Forschungszentrum für Künstliche Intelligenz, Germany
piotr.danilewski@tcs.uj.edu.pl slusallek@cg.uni-saarland.de

arXiv:1612.01325v1 [cs.PL] 5 Dec 2016

Abstract

When creating a new domain-specific language (DSL) it is common to embed it as a part of a flexible host language, rather than creating it entirely from scratch. The semantics of an embedded DSL (EDSL) is either given directly as a set of functions (shallow embedding), or an AST is constructed that is later processed (deep embedding). Typically, the deep embedding is used when the EDSL specifies domain-specific optimizations (DSO) in a form of AST transformations.

In this paper we show that deep embedding is not necessary to specify most optimizations. We define language semantics as *action functions* that are executed during parsing. These actions build incrementally a new, arbitrary complex program function.

The EDSL designer is able to specify many aspects of the semantics as a runnable code, such as variable scoping rules, custom type checking, arbitrary control flow structures, or DSO. A sufficiently powerful staging mechanism helps assembling the code from different actions, as well as evaluate the semantics in arbitrarily many stages. In the end, we obtain code that is as efficient as one written by hand.

We never create any object representation of the code. No external traversing algorithm is used to process the code. All program fragments are functions with their entire semantics embedded within the function bodies. This approach allows reusing the code between EDSL and the host language, as well as combining actions of many different EDSLs.

1. Introduction

Creating a new DSL from scratch is not an easy task. Even with the help of parser generators, such as YACC [Joh75], the developer still needs to handle several compiler-related tasks such as creating and transforming an abstract syntax tree or implementing variable lookup rules or a type system.

An alternative approach is to *embed a DSL* – define it as part of a flexible, general-purpose *host language*. The DSL is defined by a set of functions and other constructs giving the user a perception of programming in a new language.

In this paper we focus on *action code* – pieces of code that define the semantics of the language. The simplest approach for the action is to contain the intended domain program behavior directly. When the associated syntactic construct is encountered by the interpreter or parser, the domain program is executed immediately.

Alternatively, one uses some form of *staging* – a user-guided partial evaluation mechanism where the programmer annotates the code, guiding the interpreter to execute the code out of the normal order. In the context of a DSL definition, staging is used to defer the execution of the domain program. Most typical approaches of staging, such as Lightweight Modular Staging (LMS) [RO12], produce a program as an Abstract Syntax Tree (AST). The generated program, represented as an object, can be then traversed, analyzed, optimized and later compiled and run separately.

One embeds a DSL to simplify the definition creation, and to reuse code between the host language and the DSL. However, when the code is represented as an AST these traits are lost. The AST traversal and transformations must once again be defined. Moreover, the host language can no longer be seamlessly integrated with the produced program: The host code cannot be directly reused in the domain and the implementation details of the embedding can become visible to the end user [Jov+14].

Contributions

In this work we show a novel way to define language semantics without representing it as an AST object and traversing it separately. Despite that constraint, the language should still be able to produce highly efficient code.

We create *fragment functions* which conceptually act as tree nodes. Unlike existing solutions, however, their behavior is defined entirely by their body, without any external traversing algorithm. A single node can contain arbitrary code and realize a complex language logic. We show:

- How to define and connect these fragment functions.
- How to use staging to remove overhead and specify early-stage optimizations.
- How to define tasks which are typically realized through AST traversing algorithms. This includes basic operations, such as variable name lookup rules, as well as more advanced domain-specific optimizations.

Our solution has simplicity of shallowly embedded DSL and the power of deep embedding.

Embedding type	Staging method	Host language/system	Example languages created with it
standalone	structural staging	LLVM [LA04]	clang [Lat08], Kaleidoscope[1], Stacker[2]
deep embedding ↑		Lisp [Baw79], Racket [TH+11]	Pollen[3]
	functional staging	MetaML [TS97], MetaOCaml	[SBP99], QBF DSL [Cza+03]
		LMS [RO12]	[Rom+11]
⋮	dynamic staging	DeepCPS [Dan+14]	this paper
	automated staging	Impala [Lei+15]	AnyDSL[4]
shallow embedding ↓	no staging	Haskell	ImplicitCAD[5]

Figure 1: The relation between different staging methods and language creation types.

2. Related Work

2.1 Staging

We call *staging* any method that allows symbolic computation "under a lambda", that is: performing β-reduction within bodies of functions that have not been invoked yet. We define the following types of staging:

The simplest, most crude approach is *textual staging*, where a program is composed from strings containing pieces of source code. Textual staging provides no correctness guarantees about the generated program and has limited mechanisms for communication between the stages.

In *structural staging*, the program code is represented explicitly as a data structure, typically as an AST or graph. The structure can be created explicitly, for example through the LLVM instruction builders. Alternatively, the process can be hidden behind overloaded functions, operators, or templates [Vel96]. Data exchange between stages is inconvenient since it requires conversion from values to structure nodes and vice-versa. Moreover, nested staging requires a structure to describe a program building another structure, making such approach inconvenient to use.

Functional staging is a case of structural staging where the structure is represented by ordinary functions [Rey75; CKS09]. Such functions can be combined to form bigger pieces of code. The piece of code is executed by simply invoking the function representing it. In Lightweight Modular Staging (LMS) [RO12] these builder functions are hidden by overloading ordinary functions over a higher-kinded type `Rep[T]`. Consequently, the staged code does not differ much from the normal code. Unlike earlier approaches, communication is facilitated through an implicit promotion of values from `T` to `Rep[T]` where needed. Still, the structural staging is not entirely hidden and sometimes it leaks to the programmer. As a result, they must be aware of the internal usage of staging within the builders [Jov+14].

Languages such as 'C [EHK96] or MetaML [TS99] use a dedicated syntax to represent code objects. In MetaML, pieces of code can be combined, spliced, and executed. Values can be used in many stages through splice and lift operators. MetaML allows for arbitrarily deep nesting of staging. An enhanced type system is able to guarantee correctness of staging [Mog+99].

All the above techniques represent the staged code as objects within the host language. With *dynamic staging*, realized by the DeepCPS language, this view is completely abandoned [Dan+14]. There is no inherent difference between the code in different stages. Staging is controlled by first-class *staging parameters* that control the execution order by annotating instructions. Symbolic computation and deferred execution is implicitly achieved when execution occurs under a lambda.

Impala [Lei+15] has a similar view on staging. The language is less verbose and more automatic in selecting what parts of code are executed based on a single user annotation. Unfortunately, it is not easy to parametrize the staging and have it dependent on arbitrary computation.

We build our work on top of DeepCPS, making use of the language staging flexibility. We benefit from the lack of any inherent distinction between stages, as it removes the difference between the immediate action code and the program that is being built.

2.2 Language Embedding

Languages can be created as standalone products, or as a library in a flexible host language such as Haskell, ML, or Scala. The second approach is referred as *language embedding* as it defines the DSL within the *host language*.

Such an approach reduces the complexity of the language creation. Most of the compiler-specific challenges, such as having a sound type system, are already resolved in the host language and can be reused in the DSL. The author can focus only on aspects that are truly unique for the domain.

Having the common base language, DSLs can use it to relay information between the domains and the host language. For this reason, this approach is common when defining small, embedded DSLs (EDSLs). For example, Haskell has been used extensively to define geometric operations [HJ94], CAD operations[5], COM component scripting [JML98], hardware design [Bje+98], or server-side web scripting [Mei00; Thi05].

[1] Kaleidoscope: Code generation to LLVM IR, http://llvm.org/docs/tutorial/LangImpl3.html, retrieved on 12.05.2015

[2] Stacker: An Example Of Using LLVM, http://llvm.org/releases/1.1/docs/Stacker.html, retrieved on 2.12.2016

[3] Pollen: the book is a program, http://docs.racket-lang.org/pollen, retrieved on 2.12.2016

[4] AnyDSL – A Framework for Rapid Development of Domain-Specific Libraries, https://anydsl.github.io/, retrieved on 2.12.2016

[5] ImplicitCAD, http://www.implicitcad.org/, retrieved on 09.06.2016

Basic syntax:	$\lambda\,[y]\,\overline{x}.b$	`(x1, x2, ...)[y]{b}`	(lambda function)
	$\top\ \bot$ `and or not`	`always never & | !`	(staging constants, operators)
	$[e]\,v\,\overline{v}$	`@e:v v1 v2 ...`	(λ body: application)
	$[e]$ `fix` $[y]x = v$ `in` b	`@e:fix [y]x v b`	(λ body: fix-point combinator)
Syntactic sugar:	$\lambda\,[y]\,\overline{x}.\,[y]\,v\overline{v}$	`(x1, x2, ...) { v v1 v2 ... }`	(natural staging)
	$[e]\ \ (\lambda\,[y]\,x.b)\ \ v$	`@e:let [y] x v b`	(let construct)
	$[e]\,v\,\overline{v}\,(\lambda[y]x.b)$	`@e:v v1 v2 ... (x1, x2, ...)[y]b`	(last argument)
	$(\lambda[y]x.b)\,expr \xrightarrow{\beta} (\lambda[y]x.b)\,v$	`expr(x)[y] b`	(non-CPS expression $expr$)

Figure 2: Comparison between lambda calculus extended with dynamic staging (left), and the actual syntax of DeepCPS (right).

Most EDSLs focus on semantics, while inheriting the syntax of the host. The semantics is defined by pieces of code, that we refer to as *actions*. Depending on its contents we distinguish two forms of embedding:

- In *shallow embedding* the associated semantics is represented directly in the action and is executed immediately when the syntactic construct is encountered.
- In *deep embedding* a program structure representing the intended behavior is created, that can later be translated, optimized, and run separately.

This distinction is closely related to the form of staging used within the actions. Typical shallow embedding uses no staging and the code is executed together with the embedded syntax representing it. Deep embedding on the other hand uses structural staging to generate a program based on the syntax. Direct structural staging is still very common, as in Haskell/DB [LM99] or Scala [HO10]. However, more advanced staging approaches become popular in context of EDSLs as well, using LMS or MetaML [Cha+10; Ofe+13; SBP99].

This relation between the type of staging and the depth of the embedding, as depicted in Figure 1, are of particular interest to us. In our work we provide a solution which, due to the flexibility of staging, is on the verge between what can be considered shallow and deep embedding.

3. Dynamic Staging

Before we present our main work let us quickly summarize the dynamic staging and DeepCPS language as our work heavily utilizes it. The full information, including the formal definition, can be found in [Dan+14].

3.1 DeepCPS

DeepCPS is a functional language operating entirely in the Continuation Passing Style (CPS) [App07]. In Figure 2 we summarize the syntax of DeepCPS. The basic lambda value $\lambda\overline{x}.b$ is represented as `(x1,x2,...){b}`. If the lambda is the last argument in an application – a scenario common in the CPS programming – the braces surrounding the body b can be omitted. This way DeepCPS avoids the deep nesting of parenthesis, typically present in other CPS representations.

Dynamic staging is introduced through the *implicit staging parameter* `[y]` that is present in every lambda and a *staging expression* `@e:` in each lambda body. When a lambda is invoked, the implicit staging parameter is substituted by a staging constant $\top$. Staging expressions use these staging parameters to form boolean expressions. When `@e:` evaluates to $\top$, the annotated body is considered *active* and is scheduled for the execution. At the each execution step the deepest active body, containing no nested active bodies, is executed.

```
fix for(!args, startval, endval, body, endfor) {
  startval<endval (cond)
  if cond () {
    body !args startval endfor (!args2)
    startval+1 (nextval)
    for !args2 nextval endval body endfor
  } ()
  endfor !args
} in ...
```

Listing 1: A `for` loop defined as a higher-order recursive function, using the DeepCPS syntax. In each step, the user-provided `body` function is called, passing in and out a tuple of arbitrary arguments `!args`. Upon reaching the end, the `endfor` continuation is invoked.

Staging variables can be used as normal arguments, and normal variables can appear in staging expressions as well. Non-stage constants are equivalent to $\top$, while variables that are still represented only as symbolic values are $\bot$.

As syntactic sugar, the staging expression can be omitted, implicitly staging the body on the lambda it is directly contained in. When that lambda is invoked, the body is executed as if no staging was present. We refer to this as *natural staging*.

3.2 Control Flow

Control flow in DeepCPS is realized through a single built-in function `if` taking 3 arguments: `cond`, `tb`, `fb`. Depending on a boolean value `cond`, either the true branch (`tb`) or the false branch (`fb`) continuation is invoked.

Other flow functions can be defined with the help of `if`, for example as shown in Listing 1. The `for` function recursively calls itself, incrementing the initial value `startval` by 1 in each call. In each iteration, the body is invoked with the current value of `startval`. The recursion is terminated when the `endval` value is reached.

A parameter preceded by `!` accepts any excessive arguments given to a function, and packs them all into a tuple

```
... @s1: ... code 1 ... [s2] ...
...
... @s2: ... code 2 ... [s3] ...
...
... @s3: ... code 3 ... [s4] ...
```

Listing 2: The staging chain pattern. When `s1` becomes active, *code 1* is executed. Its final continuation activates `s2`, triggering the execution of the next *code* piece. This way, all the code pieces $code_n$ are executed in sequence. The code pieces can occur at arbitrary locations in the program, even in separate functions, as long as the staging values that form the chain are properly passed.

under the given name. When, for example `for` is invoked with 6 arguments, the first two are put into the tuple `args`.

The symbol `!` in front of an argument unpacks all elements of a tuple and splices them as an argument list to a function call. For convenience, all arguments that are packed or unpacked through `!`, are highlighted in italics.

3.3 Staging Chain

One of the benefits of DeepCPS that we rely on in this paper is the ability to form staging chains. A *staging chain* is a sequence of staging variables `s1`, `s2`, `s3`... where s_i stages a piece of code that is ending with a continuation defining s_{i+1}. The simplest case is shown in Listing 2, but these pieces can originate from completely independent functions. A single piece may be encapsulated in a function such as:

```
let fragment(s_n, return) {
  ... @s_n: ... code 1 ... [s_m] ...
  return s_m
}
```

4. Building Code

Consider a language designer creating a new DSL. Let us assume that the grammar for such language already exists. The precise representation of the grammar does not interest us: the grammar may be defined in a standalone parser, in a Backus-Naur form [Bac+63], or through a set of parser combinator functions, or as an embedding in another language.

The grammar contains semantic actions in a form of action functions, executed when the given grammar rule is taken. We assume that each action has means to communicate with other actions by passing arguments. We only require that the communication is performed in one direction, in the order that the actions are invoked. This can be the order of parsing, e.g. by using an L-attributed grammar [LRS74].

The question we answer in this section is how to specify these actions, so that by invoking them we generate a new program code. We use the DeepCPS syntax, but for the moment use no staging.

4.1 Immediate Execution

The most straightforward approach is to provide the intended semantic meaning of the parsed code directly in the action. Suppose for example, one implements a simple stack-based interpreter of binary arithmetic expressions as in Listing 4:

```
let push(S, v, return) { return [v, S] }
let pop(S, return) {
  S[0] (v)
  S[1] (tail)
  return v tail
}
```

Listing 3: The `push` and `pop` functions for a stack represented as a nested tuple `[v,tail]`. Each element `v` is paired with a `tail` representing the rest of the stack. After each operation, an updated stack is passed into the continuation `return`.

```
let start(return) { return [] }
let number(S, val, return) { push S val return }
let add(S, return) { pop S (arg2, S)
                     pop S (arg1, S)
                     arg1+arg2 (sum)
                     push S sum return }
let sub ... let mul ... let div ...
let end(S) { pop S $exit }
```

Listing 4: Semantic actions for an immediate interpreter of arithmetic expressions. Each action is represented as a DeepCPS function taking a stack argument `S` and returning its updated version through the continuation `return`.

- We use a stack `S`, represented as a nested tuple `[v,tail]`. For convenience, we define `push` and `pop` in Listing 3.
- Whenever a number is encountered, it is pushed onto `S` .
- When an operator is encountered, the operands are taken from `S` and evaluated, pushing the result back onto `S`.

In addition, we include the `start` action, that initializes the stack as an empty tuple, and `end`, that terminates the whole DeepCPS program through a built-in `$exit`.

Of course, the downside of this approach is that we actually never generate any code. Instead, the semantic actions compute the result immediately, as in a plain shallowly embedded DSL without any staging.

4.2 Building Code in CPS

Instead of performing the computation immediately, an action should produce a piece of a code that can be called later. We achieve that by encapsulating it within lambda functions. We introduce *builder functions* – a CPS approach for such encapsulation, consisting of 3 layers of functions:

- The *core* constitutes a piece of code that we want to generate. In CPS representation, it is a function taking any number of arguments, performing its user-defined function and then returning through the continuation with a (possibly different) set of arguments.
- The core is encapsulated by a *fragment function* which specifies the continuation for the core.
- Finally, a *builder function* acts as a constructor: When invoked with an arbitrary code as an argument, it creates a fragment function with the core containing said code.

A simple builder from Listing 5 calls its continuation `return` with a newly created fragment function. By invoking

```
let build
  (code, return) { return                          builder
    (next, return) { return                  fragment
      (!args)                                   core
        code !args (!args2) next !args2
    }
  }
```

```
let merge
  (Fprev, Fnext, return) { return
    (next, return) {
      Fnext next (core)
      Fprev core (core2)
      return core2
    }
  }
```

Listing 5: The basic builder and merge functions for our functional code building. The builder creates a fragment function that encapsulates a piece of arbitrary user code, given as an argument to the builder. The merge function connects two existing fragment functions `Fprev`, `Fnext`, forming a single new fragment containing all user code from `Fprev` and `Fnext`.

```
let number(Fprev, val, return) {
  build (S, cont) { push S val cont }
        (Fnumber)
  merge Fprev Fnumber return
}
let add(Fprev, return) {
  build (S, cont) {
          pop S (arg2, S)
          pop S (arg1, S)
          arg1+arg2 (sum)
          push S sum cont
        } (Fadd)
  merge Fprev Fadd return
}
```

Listing 6: Semantic actions for an interpreter of arithmetic expressions, creating a lambda function with the code containing all the parsed operations.

the fragment with a specified continuation `next` we create a core. The core takes arbitrary many parameters, indicated by `!`, and passes them to the user-defined `code`. Therefore, in order to form and execute a program one needs to combine:

builder(code) + fragment(next) + core(args)

Fragment functions are convenient to use because they can be combined together through a *merge function* to form a bigger fragment with the same structure. The `merge` from Listing 5 takes two fragment functions `Fprev` and `Fnext`. The core of `Fnext` is given as the continuation for the core of `Fprev`. This way the two bodies become connected, with contents of `Fnext` appearing after `Fprev`:

```
(next, return) { return                  merged fragment
  (!args)                                  merged code
    FprevCode !args (!args2)
    FnextCode !args2 (!args3)
    next !args3
}
```

The resulting fragment function has the same structure as a fragment function created by the builder, and can be further merged with other fragments to build bigger programs.

We must remark however, that the reduction described above is performed within the body of the fragment lambda returned by the `merge` function. This implies staging. Without it, we actually obtain a result that is only β-equivalent to the one shown above. Working with such fragments would incur a major overhead. We fix this problem using staging in Section 4.3.

With the help of the builders we now redefine our arithmetic expression language. In Listing 6, we no longer execute the code immediately. Instead, each of the actions creates a new fragment, and merges it with the previous one (`Fprev`).

4.3 Removing the Builder Overhead

The simple builders we presented so far are lazy in their evaluation: the merging of the fragments does not occur until the actual program is executed, incurring a potential major overhead. Instead, we would like to produce any code as if it was written directly in DeepCPS. The whole layer of builders, fragments, and core functions calls should resolve early, leaving only the user-defined code. Moreover, we would like to be able to express optimization and staging strategies in the semantic actions as well.

In this section we distinguish three execution phases:

- The *build-time* (`bt`) when the builders are being invoked.
- The *core-time* (`ct`) when all code pieces are already merged into a single fragment and are ready to be transformed into a produced program.
- The *program-time* (`pt`) when the function that we created – referred as a *program* – is invoked.

We avoid the more common terms compile-time and run-time, because this distinction can be misleading in the context of our builders and dynamic staging. This is because all the execution occurs in the very same language: DeepCPS. All code, including the programs we construct, can be executed immediately, even during build-time of another function. Thus, a program-time for one function can be a build-time for another. Such flexibility becomes lost when thinking in rigid terms of compilation and final run-time execution.

With the dynamic staging it is possible to achieve our goals. The staged builders in Listing 7 have the same structure as the original ones from Listing 5. The difference is the use of staging within the bodies of the build and merge functions.

In Listing 8 we show how the staging allows the merge to remove the overhead of having layered fragment function calls. Suppose that `merge` is invoked with two fragments `Fnext` and `Fprev` produced by the builder. As part of the invocation, the implicit staging parameter `bt` becomes $\top$ and triggers immediate execution of the code within the body of the merge fragment (Listing 8b).

```
let build
  (code, return) { return
    (next, return) { return
      (ct, pt, !args)
        @ct:   code pt !args (pt !args2)[ct]
        @next: next ct pt !args2
    }
}
```

```
let merge
  (Fprev, Fnext, return)[bt] { return
    (next, return) {
      @bt:     Fnext next (core)
               Fprev body (core2)
      @return: return core2
    }
}
```

Listing 7: The simple builders with dynamic staging. The core uses two fragment chaining patterns `ct` and `pt`. Moreover, calling `merge` triggers the `bt` stage. The fragments `Fnext` and `Fprev` are merged immediately in the body of lambda `(next, return)...`

```
build code1 (Fprev)
build code2 (Fnext)
merge Fprev Fnext (F) ...
```

(a) Step 0: Initial setup – we provide two code pieces into the builder and then try to merge the produced fragments.

```
(next, return) {                                  merge fragment
  @⊤: (next, return) { return                          Fnext
    (ct, pt, !args)
      @ct:   code2 pt !args (pt !args2)[ct]
      @next: next ct pt !args2
  }
  next (core)
  (next, return) { return                              Fprev
    (ct, pt, !args)
      @ct:   code1 pt !args (pt !args2)[ct]
      @next: next ct pt !args2
  }
    core
    (core2) @return: return core2
}
```

(b) Step 1: The `merge` function has been invoked with `Fprev` and `Fnext` arguments produced by the builder. The respective lambda values replace the symbolic names in the universal merge function.

```
(next, return) {                                  merge fragment
  @⊤: (next, return) { return                          Fprev
    (ct, pt, !args)
      @ct:   code1 pt !args (pt !args2)[ct]
      @next: next ct pt !args2
  }
    (ct, pt, !args) {                          core of Fnext
      @ct:   code2 pt !args (pt !args2)[ct]
      @next: next ct pt !args2
    }
    (core2) @return: return core2
}
```

(c) Step 2: The `Fnext` has been invoked and it has immediately returned through the `(core)...` continuation. The argument `core` is now the core of `Fnext`.

```
(next, return) {                                  merge fragment
  @return: return
    (ct, pt, !args)                        core2, core of Fprev
      @ct:   code1 pt !args (pt !args2)[ct]
                                           next, core of Fnext
      @ct:   code2 pt !args (pt !args2)[ct]
      @next: next ct pt !args2
}
```

(d) Step 3: `Fprev` has been called with the core of `Fnext` as the `next` argument. This has caused the `Fnext`'s core to be immediately invoked due to `@next:` staging.

Listing 8: Example run of staged `merge` invoked with fragment function arguments produced by the staged `build`.

The `Fnext` is given a symbolic argument `next`. The function merely invokes the returning lambda continuation, with `(ct, pt, !args){...}` as the `body` argument (Listing 8c)

This is followed by the invocation of `Fprev`. This time however, the fragment argument `next` is a concrete lambda – the previously returned core of `Fnext`. As a result, within the body of the fragment `Fprev`, the `next` function is immediately invoked. Consequently, the whole line:

```
@next:next ct pt!args2
```

is replaced with the core of `Fnext`. Only when that happens, the fragment `Fprev` returns (Listing 8d).

Suppose that we have constructed and merged all fragments representing our program. What we get is a single fragment function of the form:

```
let Fcomplete(next, return) { return
  (ct, pt, !args)
    @ct:   code1 pt !args (pt2 !args2)[ct2]
    @ct2:  code2 pt2 !args2 (pt3 !args3)[ct3]
    @ct3:  code3 pt3 !args3 (pt4 !args4)[ct4]
    ...
    @next: next ctn ptn !argsn
}
```

We can identify the use of the fragment chaining pattern from Section 3.3 twice, associated with `ct` and `pt`.

The core-time staging chain `ct` is used to stage the calls to the user functions $\texttt{code}_n$, effectively executing the contents of the core functions (hence the staging name). This resolves the process of packing and unpacking parameters into the `!args` tuples. After the core-time chain is resolved, only the code actually provided by the user remains.

If the user code has no staging (that is – has only natural staging), *all* of it is evaluated at core-time. However, by using the program-time `pt` staging chain, execution of parts of the code are delayed and spliced into `Fcomplete`. The code staged upon `pt` constitutes the program we generate.

Notice that the builders themselves never stage upon `pt` – it is merely passed around together with `!args`. In fact, `pt` can be part of `!args`, but we included it explicitly for clarity. Moreover, `!args` can include more user-defined staging variables. With these, users are able to define arbitrary additional dynamic staging within their code.

4.4 Branching via Universal Builders

The basic builders we work with so far assume that the code requires a single continuation. We say that these fragment functions are of arity 1. However, when more continuations are needed, the signature of the fragment changes and the

```
let build
  (arity, code, return)[bt] { return                    builder
    (Snext, return) {                                   fragment
      @bt: for Snext [] 0 arity
        (Snext, Tlambda, i, break, continue)[bt] {
          @Snext: pop Snext (next, Snext)
          @bt: let lambda_i
            (pt, !args2)[ct]{          continuation λ_i
              @next: next ct pt !args2
            }
          concat(Tlambda, [lambda_i]) (Tlambda)
          continue Snext Tlambda
        }
      (Snext, Tlambda) //endfor
      push Snext
        (ct, pt !args) {                                core
          @ct: code pt !args !Tlambda
        } (S)
      @return: return S
  } }
let merge
  (Fprev, Fnext, return)[bt] { return
    (Snext1, return) {
      @bt:      Fnext Snext1 (Snext2)
                Fprev Snext2 (Snext3)
      @return: return Snext3
    }
  }
```

Listing 9: The universal builders handling fragments of arbitrary arity. The fragments exchange a stack of continuations `Snext`, rather than a single continuation `next`. In the `for` loop, the fragment takes as many elements from the stack as needed, defined by the `arity` parameter, and pass them into the user-defined `code` as continuations.

builders have to be adjusted. For example, when encapsulating a condition function `if` that takes two continuations:

```
build2 (trueCont falseCont) {
         if cond trueCont falseCont
       }
       (F) ...
```

one needs to use a version of the builder which creates a fragment function of arity 2, for example defined as `build2`:

```
let build2
  (code, return) { return
    (next1, next2, return) { return
      (!args)
        code !args (!args2) { next1 !args2 }
                   (!args3) { next2 !args3 }
  } }
```

Fragment functions of different arity present a problem as they cannot be directly used with the `merge` function presented earlier. On the other hand, creating a merge function to match fragments of different arity would be ineffective – the user would have to use several versions of the function, where the distinction could be automated.

Instead of using specialized builder and merge function, in Listing 9 we introduce a universal solution for any arity.

The main difference is that we no longer pass a single `next` continuation into the fragment functions. Instead, we pass a stack `Snext`, containing arbitrary many continuations. We use the same stack implementation as previously in Listing 3.

```
(Snext, return) {
  @Snext: pop Snext (next, Snext2)
  @Snext2: pop Snext (next2, Snext3)
  @return: return
  [(ct, pt, !args) {          core (top stack element)
    @ct: code pt !args ![
      (pt, !args) { @next: next ct pt !args },    λ_0
      (pt, !args) { @next2: next2 ct pt !args }] λ_1
  }, Snext3]
}
```

Listing 10: A fragment function obtained by invoking the universal `build` with arity 2.

When creating a new fragment through the builder, the user specifies the intended `arity`. This specifies the number of continuations passed by the builder to the user-provided `code` that follow the standard `!args`. These continuations are kept in the `Tlambda` tuple, created at build-time in the `for` loop at the beginning of the builder's fragment function. Each of the `next` values used by these continuations is obtained from the `Snext` stack as soon as it is available.

The universal merge function does not differ much from its simple predecessor. However, this time we do not pass a single core argument between the fragments, but the whole continuation stack `Snext`. The code within the universal builder, through the use of `lambda_i` performs the actual merging. The `merge` function merely specifies which two fragment functions are connected.

Note that the `push` operation on the stack is performed in build-time even if `Snext` is a symbolic value. This allows us to create stack with top element known, and unknown remaining tail. This can be seen for example in Listing 10. Such element can be later used in another fragment to resolve the merge, even if the whole stack is not provided.

The logic of the construction through these builders is from back to front: when the continuations representing the ending part of the program are known, they are put on the stack and can be merged with a new fragment of code preceding it. Despite that it is not necessary to call these functions in that one particular order. Thanks to the operations on incomplete stack, the `Fnext` can be successfully merged to `Fprev` even if the former is just a short piece of code without any known continuations.

5. Examples

In the previous section we have shown how arbitrary code can be created using a functional approach. We have explained how dynamic staging removes the overhead of the layered construction. Let us now show how the builders can be used in practice.

5.1 Arithmetic Expressions

At the beginning of Section 4, in Listing 4 we provide semantic actions for parsing arithmetic expressions. In there, however, all the computations are performed immediately during parsing. With the help of the builders, the complete set of semantic actions would look like in Listing 11.

```
let start(return) {
  build 1
    (pt, exit, cont) { cont pt exit [] }
    return
}
let number(Fprev, val, return) {
  build 1 (pt, exit, S, cont) {
            push S val (S)
            cont pt exit S
          } (F)
  merge Fprev F return
  }
let add(Fprev, return) {
  build 1 (pt, exit, S, cont) {
            pop S (arg2, S)
            pop S (arg1, S)[ct]
          @pt: arg1+arg2 (sum)[pt]
          @ct: push S sum (S)
            cont pt exit S
          } (F)
  merge Fpref F return
}
let sub... let mul... let div...
let end(Fprev, return) {
  build 0 (pt, exit, S) {
            pop S (v, S)
          @pt: exit v
          } (F)
  merge Fpref F (Fcomplete)
  Fcomplete [] (program)
  return (exit)[pt] {
           @program: program pt exit
         }
}
```

Listing 11: Semantic actions for an interpreter of arithmetic expressions. Each action builds a fragment with the provided user code. The user code employs a number stack to accumulate intermediate values. The ending action "seals" the chain with a fragment of arity 0, taking no more constructed continuations.

The `start` action is executed first, at the beginning. Apart from the continuation, it takes the function-time staging variable, and the `exit` continuation that should end the program we are building. Moreover, we initialize the number stack with an empty value `[]`.

The `number` and `add` actions construct a fragment of arity 1, similarly to what was done in Listing 6. All operations on the number stack are performed at core-time, even if the numeric values, such as `sum` are not known at that time. Only the actual addition is delayed to be part of the program we generate, executing at program-time.

Finally, in the `end` semantic action we seal the fragments. That means, we provide the last fragment with arity 0, which calls `exit` to end the program. The fragment representing the complete program is invoked with an empty stack as `Snext`, returning a `program` which is a naked user code without any encapsulation. By invoking the program under the lambda `(exit)[pt]...` we trigger the core-time stage, leaving only the code staged upon `pt` in the body of that lambda.

The presented approach using a number stack is not the only way to realize an arithmetic expression semantics.

```
let start(return) {
  build 1 (pt, exit, cont) {
            cont pt (pt, arg)
            exit arg
          } return
}
let number(Fparent, val, return) {
  build 0 (pt, parent) { parent pt val }
    (F)
  merge Fparent F return
  }
let add(Fparent, return) {
  build 2 (pt, parent, left, right) {
                left pt (pt, arg1)
                right pt (pt, arg2)
            @pt: arg1+arg2 (sum)[ft]
            @ct: parent pt sum
          } (F)
  merge Fparent F return
}
let sub... let mul... let div...
let end(Fparent, return) {
  Fparent [] (program)
  return (exit)[pt] {
           @program: program pt exit
         }
}
```

Listing 12: Semantic actions for the same interpreter of arithmetic expressions. Here, however, the fragments are connected in an AST-like fashion. Operators are inner nodes of arity 2, while numbers are leafs with arity 0.

Alternatively, the fragments can be connected in an AST-like fashion as in Listing 12. The action associated with the operator creates a fragment function of arity 2. Each child fragment represents a subexpression on the left and right side of the operator, respectfully. A literal number, on the other hand, acts as a leaf in the tree of fragments and has arity 0.

In both approaches, we rely on the language grammar to invoke the semantic actions in the correct order. In particular, the grammar handles the precedence of the operators.

Regardless of whether we use a number stack, or AST-like fragments, we obtain the same end result. This is because the merging and calls into each user code function is resolved early. Only instructions in the `pt` chain remain. They are the same in both approaches. If, for example, we try to parse `1+4*2+3`, the produced program for both cases is exactly:

```
(exit)[pt] {
  @pt: 4*2 (prod)[pt]
  @pt: 1+prod (sum)[pt]
  @pt: sum+3 (sum)[pt]
  @pt: exit sum
}
```

5.2 Custom Environment

Suppose that a DSL requires name binding. With functional code building we do not have to rely on a single approach or follow the rules of the host language. Name binding can be handled explicitly by the user code, possibly in an early stage to avoid program-time overhead.

In order to manage the names, we need a map M that maps string names to values. We do not care how M is

```
let reference(Fparent, name, return) {
  build 0 (pt, M, parent) {
            M.lookup(name) (value)
            parent pt valul
          } (F)
  merge Fparent F return
}
let bind(Fparent, name, return) {
  build 1 (pt, M, parent, child) {
            child pt (pt, arg)
            M.bind(name, arg) ()
            parent pt arg
          } (F)
  merge Fparent F return
}
```

Listing 13: Extending the AST-based interpreter of arithmetic expressions to support name binding. The parameter `M`, acting as a binding map, is passed through all user code functions. The insertion and lookup is performed at build-time, leaving no trace of it at function-time.

implemented. It may be implemented in a purely functional manner, with each operation producing a new M value. It may use stateful memory operations. It may even use external C functions for efficiency purposes. It may also feature a mechanism to reference other maps to support nested scopes. Since dynamic staging is explicit and does not rely on program analysis, any implementation can be used with our approach.

We only require that `M.insert(name, value)` should bind, possibly symbolic, `value` to `name` in the map M, while `M.lookup(name)` should return the previously stored `value`.

As an example, we can now augment our arithmetic expression language to support name binding, as in Listing 13. All map operations are performed at core-time, allowing the map itself to be removed before the program-time.

This is the simplest example of an *auxiliary computation* that can be added into our functional code builders. Such computation does not constitute the main program logic. Instead, it is an additional support code that can be used for management, verification or domain-specific optimization.

5.3 Domain-Specific Optimization

Typically, an optimization is specified as a separate pass over the program structure. In such a pass, characteristic patterns are found and the AST is transformed, often violating the inherent node semantics.

With the functional builders and dynamic staging the optimization can often be specified as a function as well. Consider for example our arithmetic expression language operating on matrices of arbitrary dimensions. Each matrix `M` is represented by its dimensions `M.dim` known statically, and its contents `M.data` that are known only at program-time.

Consider a problem where a sequence of matrices has to be multiplied. Finding the optimal sequence of multiplication is a classic problem that can be solved through dynamic programming [Cor+09, Chap. 15.2]. Suppose that the algorithm is implemented in a function `matrix_chain_order(d)`,

```
let mulStart(Fparent, return) {
  build 1 (pt, parent, children) {
            children pt [] parent
          } (F)
  merge Fparent F return
}
let mulNext(Fparent, return) {
  build 2 (pt, Ms, parent, child, continue) {
             child pt (pt, M)
             concat(Ms, [M]) (Ms)
             continue pt Ms parent
           } (F)
  merge Fparent F return
}
let mulEnd(Fparent, return) {
  build 0 (pt, Ms, parent) {
    count(Ms)-1 (opCount)
    get M.dim from all in Ms (dims)
    matrix_chain_order(dims) (order)
    for pt Ms 0 opCount
      (pt, Ms, i, break, continue) {
           order[i] (idx)
           Ms[idx].data (M1)
           Ms[idx+1].data (M2)[ct]
      @pt: M1*M2 (prod)[pt]
      @ct: [Ms[idx+1].dim.x,Ms[idx].dim.y] (dim)
           Ms.delete(idx..idx+1) (Ms)
           Ms.insert(idx,[dim,prod]) (Ms)
           continue pt Ms
      } (pt, Ms)
    Ms[0] (finalProd)
    parent pt finalProd
  } (F)
  merge Fparent F return
}
```

Listing 14: Chain matrix multiplication as an example of domain-specific optimization. The `mulNext` action collects all the matrices, but the classic algorithm is contained entirely within the `mulEnd` action. There, the optimal order is found in `matrix_chain_order`. Based on the result, matrices are multiplied in a `for` loop. All the optimization is performed in the core-time staging chain. At program-time the loop is unrolled and only the multiplication operations remain, referencing directly the right matrices.

which for given tuple d of matrix dimensions computes a sequence describing the optimal multiplication order. As before, the implementation of such function does not concern us.

In Listing 14 we show how such optimization can be specified in the semantic actions. The `mulStart` is invoked once at the beginning of the sequence, initializing the list of matrices `Ms` with an empty tuple. The `mulNext` is used for every member of the matrix chain and accepts 2 continuations. The `child` represents a matrix subexpression yielding a single result that needs to be multiplied. The `continue` refers to the remainder of the matrix chain.

Finally, at the end, in `mulEnd` we have all the data to launch the algorithm. The dimensionality information is extracted from `Ms` and the optimal multiplication order is found. With each iteration of the `for` loop a single line of code is produced, staged for program-time, which multiplies

the matrices selected by the optimal order. All other support code: launching the algorithm, managing the matrix tuple and computing the dimensions of the result are executed early, at core-time.

6. Discussion

Let us discuss how our functional approach for code building relates to both the shallow and the deep language embeddings.

Having the semantic action, containing the intended program logic represented as a function is typical for shallow embedding. Every host language function can be used directly in the DSL and vice-versa, any program constructed by the DSL can be used in the host. This is in contrast to deep embedding where host programs must be converted to an AST representation. Such conversion is difficult to achieve transparently especially when the built-in functions and constructs, such as conditionals, are encountered.

Typical shallowly embedded DSLs have an overhead stemming from the additional layers of functions tied to the language syntax. As we have shown in Section 4.3 however, dynamic staging can resolve all overhead coming from the embedding, producing only the intended user code. Moreover, unlike typical shallow embeddings, the produced code is not directly tied to the DSL syntax. The actual computation may occur later or even in a different order, as in the chained matrix example in Section 5.3. It is also possible for a DSL to produce not one but two or more chains of fragment functions. Such technique can be used to generate a multi-pass interpretation in a single parsing, e.g. where circular definitions are present.

On the other hand, our functional builders resemble the deep embedding as well: Actions produce code objects and the code is built incrementally, by connecting the fragment functions in a tree structure similar to an AST.

Nodes in our graph structure are self-descriptive, i.e. their behaviour is entirely defined within the functions they contain. As shown in Section 5.2 these functions may contain much more than just the core semantics. The user may define many staging phases and incorporate additional, auxiliary computation within it. It may be used to define custom name lookups, but also as an analysis step, optimization, or even perform custom type checking [Her10].

Unlike the deep embedding however, we do not support tree transformations. Optimization is not achieved through analysis of the program graph, but through the partial evaluation enabled by dynamic staging. We do not claim that it is the only way programs should be optimized: careful analysis of the code is needed to generate fast, low-level machine code. What we claim, instead, that such low-level techniques are usually not needed at higher-level DSL definitions and generally should be avoided because of their complexity.

This paper focuses on the conceptual aspects of building code through dynamic staging. However, it is important to state that the presented approach has been fully implemented in the DeepCPS interpreter and compiler. This includes the examples presented here, as well as more complex DSLs, such as a language for grammar specification and for stencil code generation. In all the cases we were able to specify domain optimization, the interpreter was able to remove any residual overhead and the compiler produced an efficient machine code.

7. Conclusions and Future Work

We have presented a novel approach to code building. We use an entirely functional approach, where each builder produces a fragment function containing an arbitrary program code. We have shown how builders are defined and argued how dynamic staging allows us to remove the overhead coming from such encapsulation.

In Section 5 we have given simple examples of how typical language design problems can be handled using builders and dynamic staging. This includes domain-specific optimizations, which are also represented as additional code. It is more natural for for a domain-expert to implement an optimizing algorithm as a program, rather than as an AST transformation. The whole optimization algorithm can be hidden away in a single builder node, or even provided as a separate function, independently of the DSL definition.

We did not include the performance numbers for the machine code produced by our DSLs as we do not evaluate the efficiency of the DeepCPS compiler. Such evaluation has been done in [Dan+14]. Note, that we are able to produce the DeepCPS code containing no unnecessary residual operations originating from the builders. The back-end compiler receives the same input as if the code was written without the builders, and produces the same result.

Future Work

The focus in this paper is only on the semantic actions of a language. For a fully functioning DSL this is not enough – a parser or an embedding is needed. We did not provide any integration of our semantic actions with a parser, nor explained how to embed a DSL in DeepCPS. Such integration is in our current focus and we have made substantial advancements in that area. This is however not in the scope of this paper.

In Section 5.2 we described how dynamic staging can be used to define an auxiliary computation. The same approach can be used to build custom type systems. In theory, any type system can be defined as a staged computation and used in a DSL. To our knowledge this possibility has not yet been fully explored and requires further research. This may be particularly interesting in the context of DeepCPS as dynamic staging may bring dynamic and static typing closer together, in a similar way as it did for shallow and deep embeddings.

We also hope to explore other uses for functional code building, not directly related to language definitions. As each fragment can be an arbitrary large piece of code, the same technique can be used to build programs at bigger granularity. For example, this technique may become viable for data-flow programming, and semi-automatic, profiled assembly.

References

[App07] Andrew W. Appel. *Compiling with Continuations*. Cambridge University Press, 2007.

[Bac+63] J. W. Backus et al. "Revised Report on the Algorithm Language ALGOL 60". In: *Commun. of the ACM* 6.1 (Jan. 1963). Ed. by P. Naur, pp. 1–17.

[Baw79] Alan Bawden. *Quasiquotation in Lisp*. Artificial Intelligence Laboratory, Standford University, Feb. 1979.

[Bje+98] Per Bjesse et al. "Lava: Hardware Design in Haskell". In: *ICFP*. 1998, pp. 174–184.

[Cha+10] Hassan Chafi et al. "Language Virtualization for Heterogeneous Parallel Computing". In: *OOPSLA*. 2010, pp. 835–847.

[CKS09] Jacques Carette, Oleg Kiselyov, and Chung-chieh Shan. "Finally Tagless, Partially Evaluated: Tagless Staged Interpreters for Simpler Typed Languages". In: *Journal of Functional Programming* 19.5 (Sept. 2009), pp. 509–543.

[Cor+09] Thomas H. Cormen et al. *Introduction to Algorithms, Third Edition*. The MIT Press, 2009.

[Cza+03] Krzysztof Czarnecki et al. "DSL Implementation in MetaOCaml, Template Haskell, and C++". In: *Domain-Specific Program Generation*. 2003, pp. 51–72.

[Dan+14] Piotr Danilewski et al. "Specialization through Dynamic Staging". In: *GPCE*. 2014.

[EHK96] Dawson R Engler, Wilson C Hsieh, and M Frans Kaashoek. "'C: A Language for High-level, Efficient, and Machine-independent Dynamic Code Generation". In: *POPL*. 1996, pp. 131–144.

[Her10] Stephan Andreas Herhut. "Auxiliary Computations — A Framework for a Step-Wise, Non-Disruptive Introduction of Static Guarantees to Untyped Programs Using Partial Evaluation Techniques". PhD thesis. University of Hertfordshire, 2010.

[HJ94] Paul Hudak and Mark Jones. *Haskell vs. Ada vs. C++ vs. Awk vs. ... An Experiment in Software Prototyping Productivity*. Tech. rep. Department of Computer Science, Yale University, 1994.

[HO10] Christian Hofer and Klaus Ostermann. "Modular Domain-specific Language Components in Scala". In: *GPCE*. 2010, pp. 83–92.

[JML98] Simon P. Jones, E. Meijer, and D. Leijen. "Scripting COM Components in Haskell". In: *ICSR*. 1998.

[Joh75] Stephen C. Johnson. *YACC – Yet Another Compiler-Compiler*. Tech. rep. Bell Laboratories, 1975.

[Jov+14] Vojin Jovanovic et al. "Yin-Yang: Concealing the Deep Embedding of DSLs". In: *GPCE*. 2014, pp. 73–82.

[LA04] Chris Lattner and Vikram Adve. "LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation". In: *CGO*. 2004, pp. 75–86.

[Lat08] Chris Lattner. "LLVM and Clang: Next generation compiler technology". In: *The BSD Conference*. 2008.

[Lei+15] Roland Leißa et al. "Shallow Embedding of DSLs via Online Partial Evaluation". In: *Proceedings of the 14th International Conference on Generative Programming: Concepts & Experiences (GPCE)*. ACM. Pittsburgh, PA, USA, 2015, pp. 11–20.

[LM99] Daan Leijen and Erik Meijer. "Domain Specific Embedded Compilers". In: *DSL*. 1999, pp. 109–122.

[LRS74] P.M. Lewis, D.J. Rosenkrantz, and R.E. Stearns. "Attributed translations". In: *Journal of Computer and System Sciences* 9.3 (1974), pp. 279 –307.

[Mei00] Erik Meijer. "Server Side Web Scripting in Haskell". In: *Journal of Functional Programming* 10.1 (Jan. 2000), pp. 1–18.

[Mog+99] Eugenio Moggi et al. "An Idealized MetaML: Simpler, and More Expressive." In: Lecture Notes in Computer Science. Springer, 1999, pp. 193–207.

[Ofe+13] Georg Ofenbeck et al. "Spiral in Scala: Towards the Systematic Construction of Generators for Performance Libraries". In: *GPCE*. 2013, pp. 125–134.

[Rey75] John Reynolds. "User-defined Types and Procedural Data Structures as Complementary Approaches to Data Abstraction." In: *New Directions in Algorithmic Languages*. 1975.

[RO12] Tiark Rompf and Martin Odersky. "Lightweight Modular Staging: A Pragmatic Approach to Runtime Code Generation and Compiled DSLs". In: *Communications of the ACM* 55.6 (June 2012), pp. 121–130.

[Rom+11] Tiark Rompf et al. "Building-Blocks for Performance Oriented DSLs". In: *Proceedings IFIP Working Conference on Domain-Specific Languages*. 2011, pp. 93–117.

[SBP99] Tim Sheard, Zine el-abidine Benaissa, and Emir Pasalic. "DSL Implementation Using Staging and Monads". In: *DSL*. 1999, pp. 81–94.

[TH+11] Sam Tobin-Hochstadt et al. "Languages as Libraries". In: *PLDI*. 2011, pp. 132–141.

[Thi05] Peter Thiemann. "An Embedded Domain-specific Language for Type-safe Server-side Web Scripting". In: *ACM Transactions on Internet Technology* 5.1 (Feb. 2005), pp. 1–46.

[TS97] Walid Taha and Tim Sheard. "Multi-Stage Programming with Explicit Annotations". In: *ACM SIGPLAN Notices* 32.12 (1997), pp. 203–217.

[TS99] Walid Taha and Tim Sheard. "MetaML and Multi-Stage Programming with Explicit Annotations". In: *Theoretical Computer Science*. 1999.

[Vel96] Todd Veldhuizen. "C++ Gems". In: ed. by Stanley B. Lippman. SIGS Publications, Inc., 1996. Chap. Expression Templates, pp. 475–487.